



The GRID as a Platform for Communication, Collaboration and E-Science

Lawrence Wilcock and Steve Hinde
Hewlett Packard Laboratories, Bristol, UK
lw@hplb.hpl.hp.com

Outline

- Important trends in communications
 - Technical & social
- Next generation communication applications
 - Characteristics and requirements
- Grid-based Collaboration Platform
 - Vision, approach, and assumptions
- Requirements for Communication Spaces
 - Control & management of virtual meeting places
- Media Grid
 - Global, media-processing fabric for real-time interactive communications

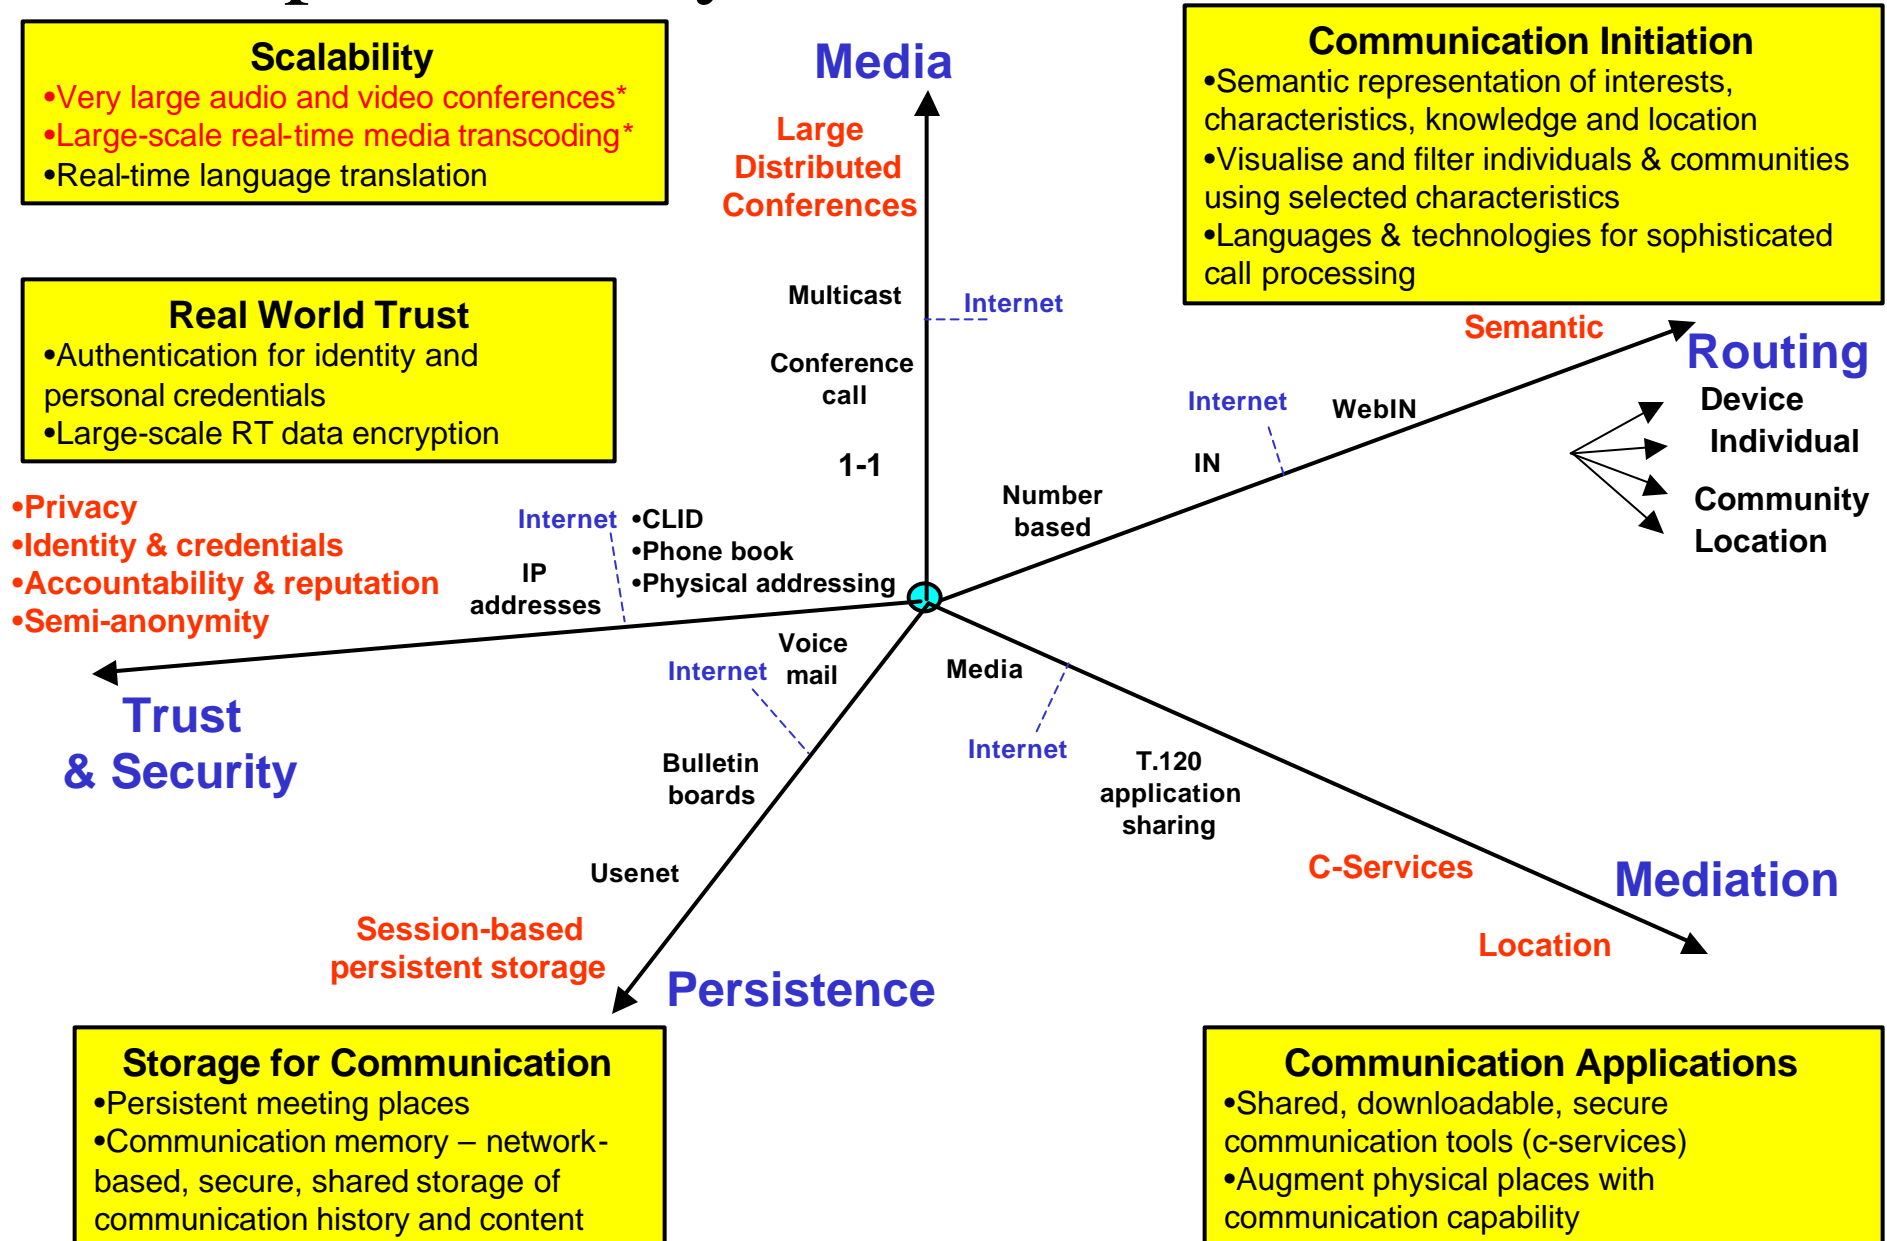
Communications in E-science

- Almost everything we do involves communication with others
- Increasing trend towards collaborative research
 - Large multi-disciplinary teams
 - Span organisational and geographic boundaries
 - Faster results, more ambitious goals
- Examples
 - Virtual laboratories
 - Online Conferences
 - E-learning and Science Lectures
- Where appropriate leverage existing technologies & commercial infrastructure
- Culture of innovation & experimentation
 - Create new technologies and paradigms when needed
 - Web, very high-speed Internet networks, GRID
 - Results feedback into Industry

Important trends in communication

- Technological
 - Greater use of computation and storage in the Network to facilitate communication
 - Exponential increase in bandwidth
 - Richer media types => greater demand for computation
 - Heterogeneity of communication devices
 - Telephone, PDA, games console, DSL-connected PC, powerful workstation on multi-megabit Internet 2
 - Lowest common denominator everywhere not acceptable
- Social
 - Increased use of Internet to communicate and collaborate in groups, teams and online communities
 - People meet in much larger groups
 - Create demand for large-scale deployment of multi-party communication applications
 - Reproduce real-world interactivity and productivity
 - Requires new paradigms and supporting technologies

Computationally Enhanced Communications



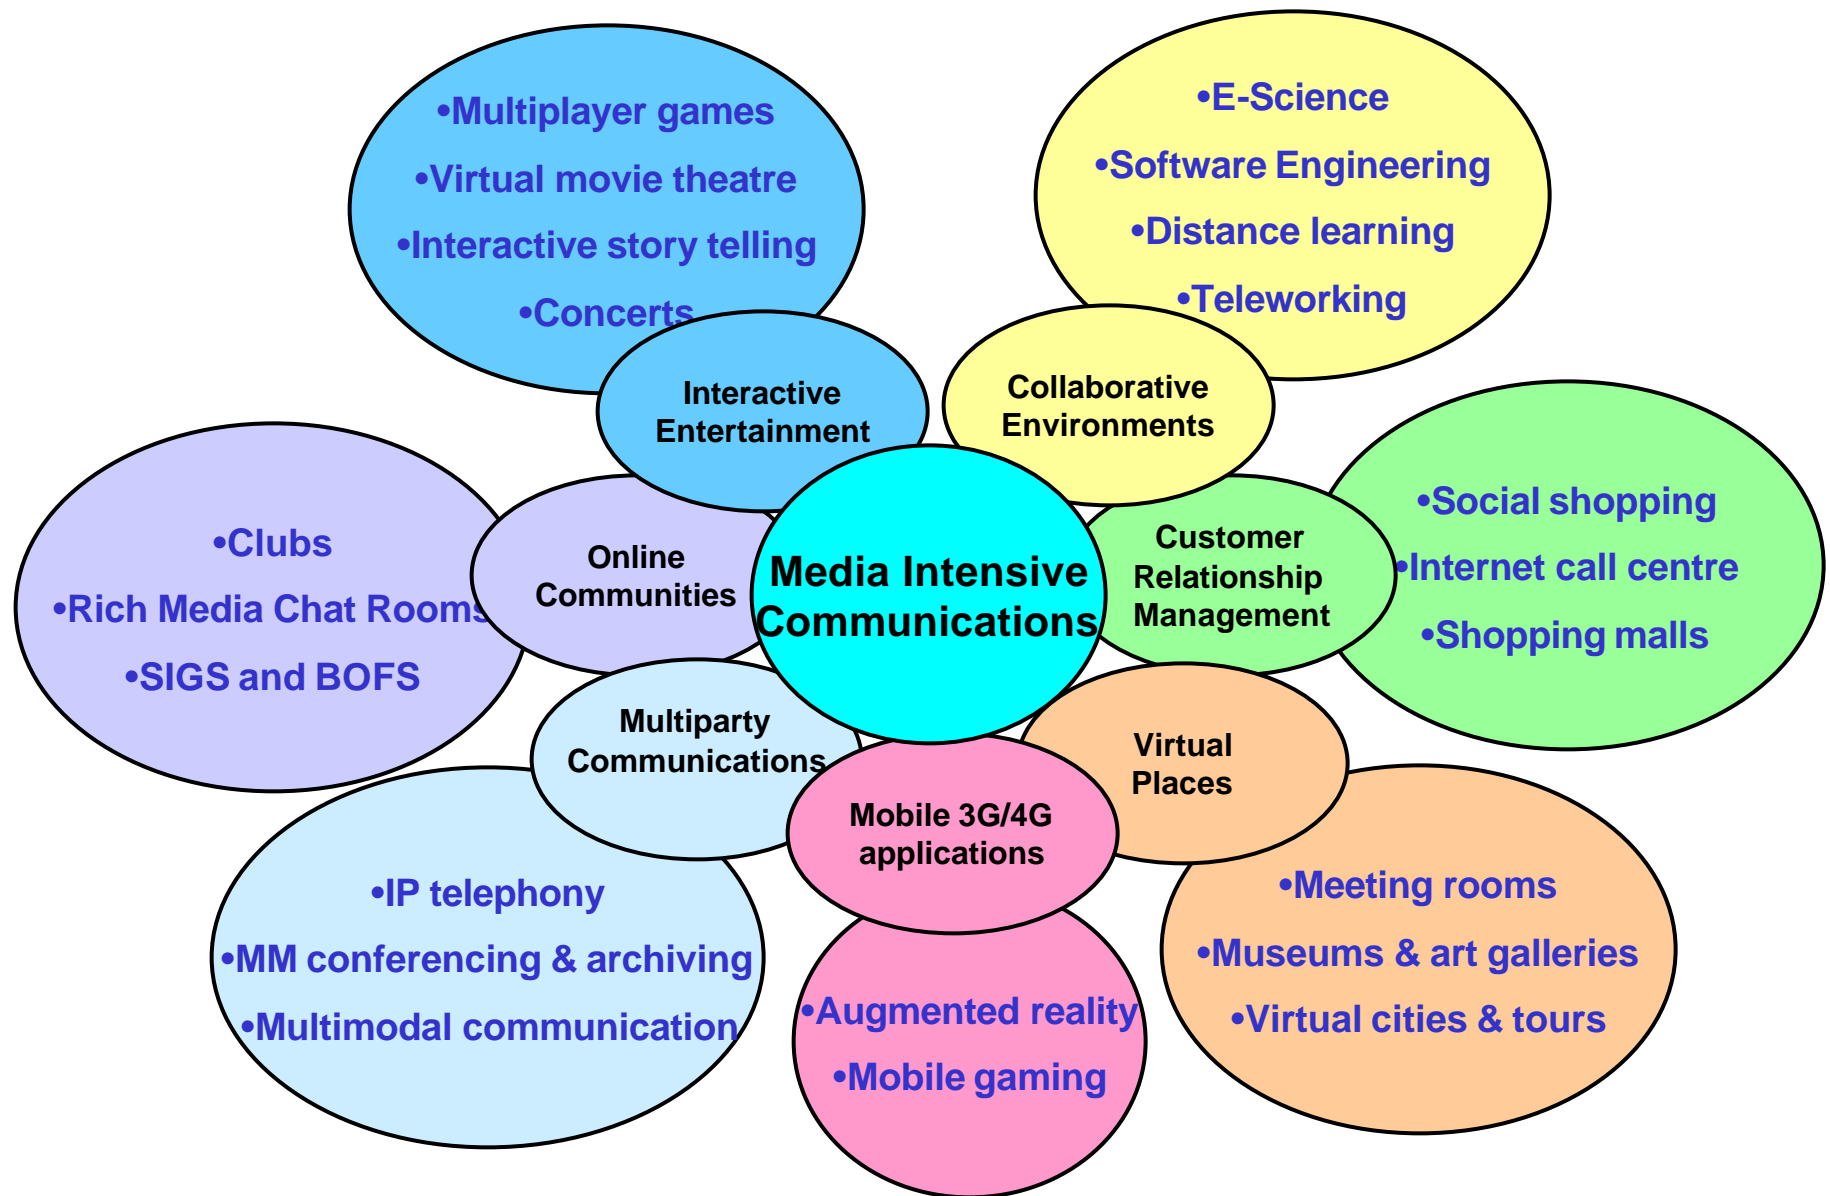
Computation Resource Crisis

- Even simple operations aggregate to huge computation demand
 - billions of users
- Particularly for real-time media
- Telephone network has insufficient compute capacity for audio mixing
 - Provisioned on assumption of 2-party calls
 - Cannot support vision of “Every call is a conference call”
- Problem:
 - Where will the compute and storage resources come from?
 - How will they be managed?
- Answer: the GRID
 - Central Office of the Future
 - Important to understand requirements for the GRID
 - Real-time media, 24x7 availability, security, QoS

Outline

- Important trends in communications
 - Use of computation
- ➔ • Next generation communication applications
 - Characteristics and requirements
- Grid-based Collaboration Platform
 - Assumptions, vision, and approach
- Requirements for Communication Spaces
 - Communications platform for constructing virtual meeting places
- Media Grid
 - Global, GRID-based media processing fabric for real-time interactive communications

Large-scale Multi-Party Communications



Characteristics of Next Generation Communications Apps

- Structure similar to real world
 - modelled as virtual meeting places
 - Environment: region, building, room, zone, group
 - have physical properties - spatial extent, acoustics
 - notion of position is important
 - Leads to usability scalability for large numbers of participants
 - Multiple simultaneous speakers
 - People can move around to naturally form groups
- Persistent
 - meeting places exist independently of participants
- Huge range in scale
 - From 2 people (& an automated agent)
 - to thousands in online game or conference


The Net as a Place

The Net as a Page




Information Delivery

- Single-user browsing
- Magazine pages
- Information searching
- 1-1 Streamed media



The Net as a Place



- Communication
- Interaction
- Real-time social media
- Online communities
- Interest groups
- Collaborative science

Communication GRID

Implications for Media Processing

- Huge demand of computation, storage, and bandwidth
 - Mixing, transcoding, transcription, voice obfuscation, TTS, ASR, media streaming & recording
- Rich media pushing computational demands even further
 - Spatial audio, real-time 3D video avatars
 - Unique mix required for each participant
 - mix contains members of own group & connected groups
- Heterogeneity => Network adaptation for diverse edge devices
 - Adapt to processing and bandwidth capabilities
 - Trade off computation for bandwidth
- Scalability
 - billions of small conferences, and massive online environments

Outline

- Important trends in communications
 - Use of computation
- Next generation communication applications
 - Characteristics and requirements
- ➔ • Grid-based Collaboration Platform
 - Assumptions, vision, and approach
- Requirements for Communication Spaces
 - Communications platform for constructing virtual meeting places
- Media Grid
 - Global, GRID-based media processing fabric for real-time interactive communications

Implications for the Infrastructure

Requirements

Communication Spaces

Session Management

- *Construct virtual meeting places*
- *Support metaphor of real-world structure*
- *Describe interconnectivity relationships between groups*

Control Plane

Media Grid

Media Processing

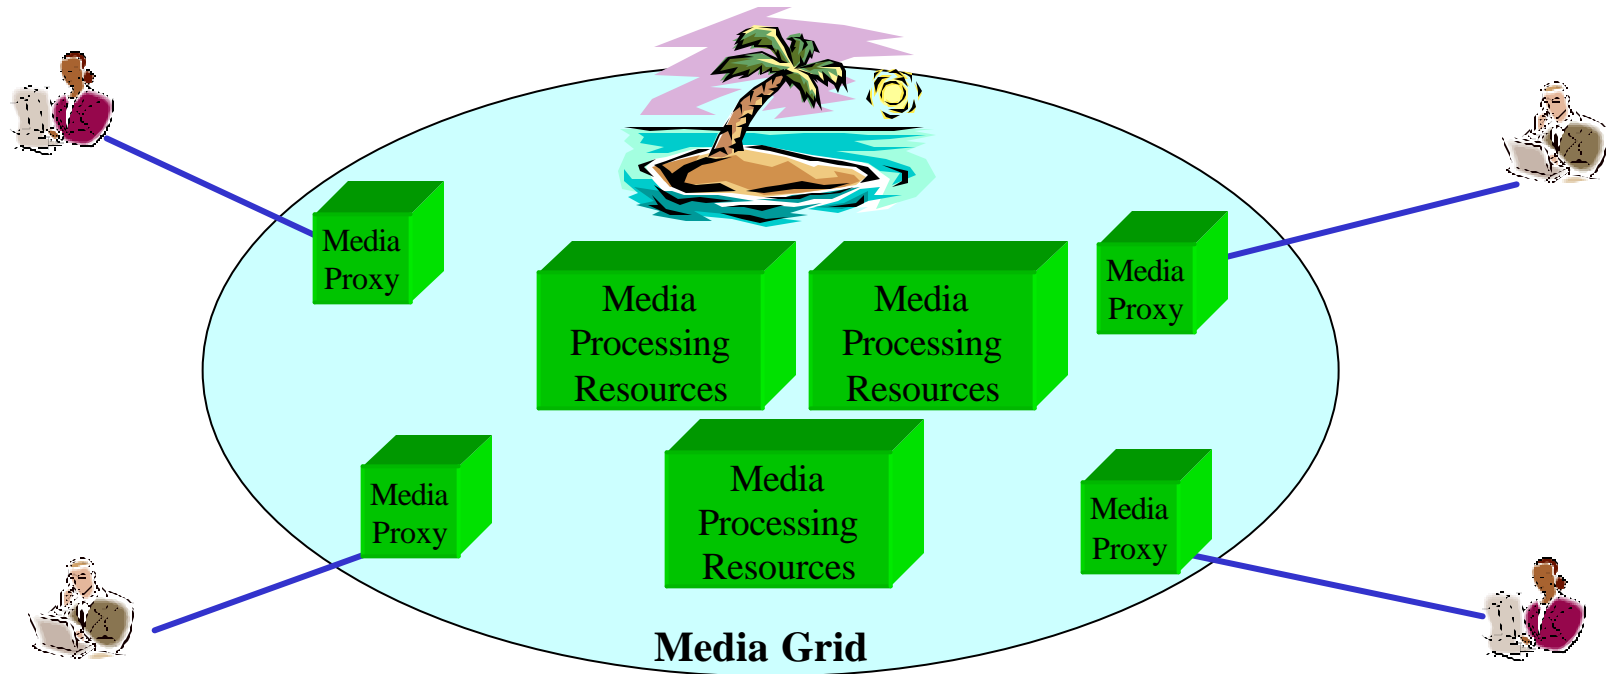
- *Scalable media processing*
- *Mixing, transcoding, network & device adaptation*
- *Complex topologies of media resources*
- *Support new media types*

Media Plane

Vision

*Enable next-generation communications applications by
creation of
scalable, planetary-scale Media Grid and Control architecture,
composed of mesh of network-resident media-processing resources,
allocated from a global pool,
And managed using GRID technologies*

Media Grid for Communications



- Global, distributed pool of media processing resources
- Connected by dedicated well-provisioned IP network
- Media resources are composable into complex meshes
- Media Proxy is user “on ramp” to Media Grid
 - mixing, transcoding, and device adaptation
- Media Resources built on GRID technologies
 - Resource description and allocation, resource virtualisation, service interfaces, security

Why Network-Hosted Media Processing?

- Edge solutions do not scale well
 - Not all end devices have required processing and bandwidth capability
 - Adapt in network – stream merging, compression, divide and conquer
- Multicast & real-time capable networks are not ubiquitous
 - do not reach all endpoints
- However, can control real-time properties within the Network
- Allows statistical sharing of expensive specialised resources
- Facilitates seamless device mobility


Assumptions

- Scalability limited by demand for media processing and bandwidth
- Real-time processing and storage of media can now be performed on a large scale in software on general purpose computing platforms
 - HP OpenCall Media Platform – 2000 IVR channels on UNIX SMP

Outline


- Important trends in communications
 - Use of computation
- Large-scale multi-party communications
 - Characteristics and requirements
- Grid-based Collaboration Platform
 - Assumptions, vision, and approach
- ➔ • Requirements for Communication Spaces
 - Communications platform for constructing virtual meeting places
- Media Grid
 - Global, GRID-based media processing fabric for real-time interactive communications

Spaces - simple meeting place




- Similar concerns to a Session
 - Participants communicate in a Space
 - Spaces control membership
 - Invitation & join mechanism
- Extend with VE concepts
 - Environment tiled with Spaces
 - Can have spatial extent
 - Can be connected together
 - Define audio physics
 - Acoustics, reverberation, echo

Spaces – company broadcast or conference



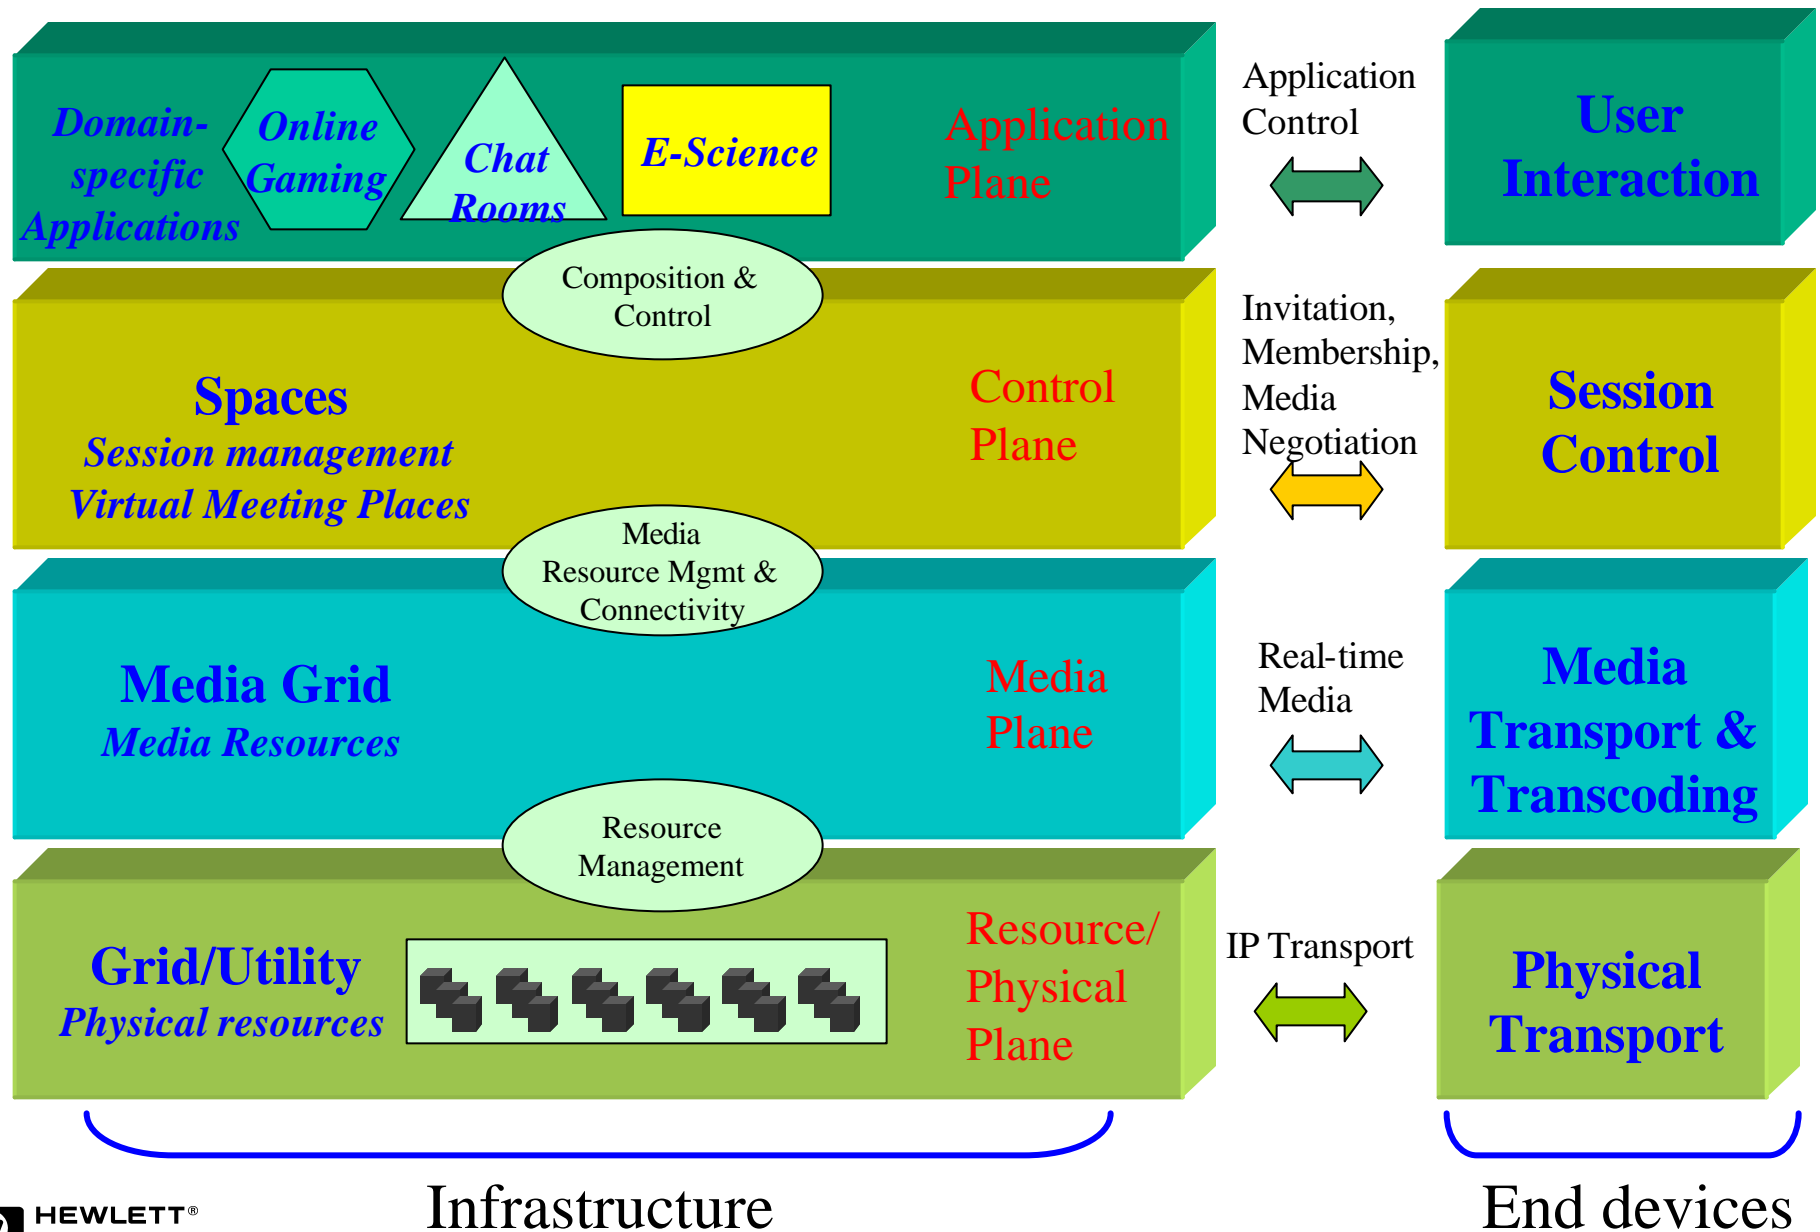
- Create complex media relationships between groups
 - Control flow of media
 - Some listening, some speaking
 - Private conversations

Spaces – pushing scalability



- Millions of Spaces, billions of users
- Massive online environments
- Surfable Web of places

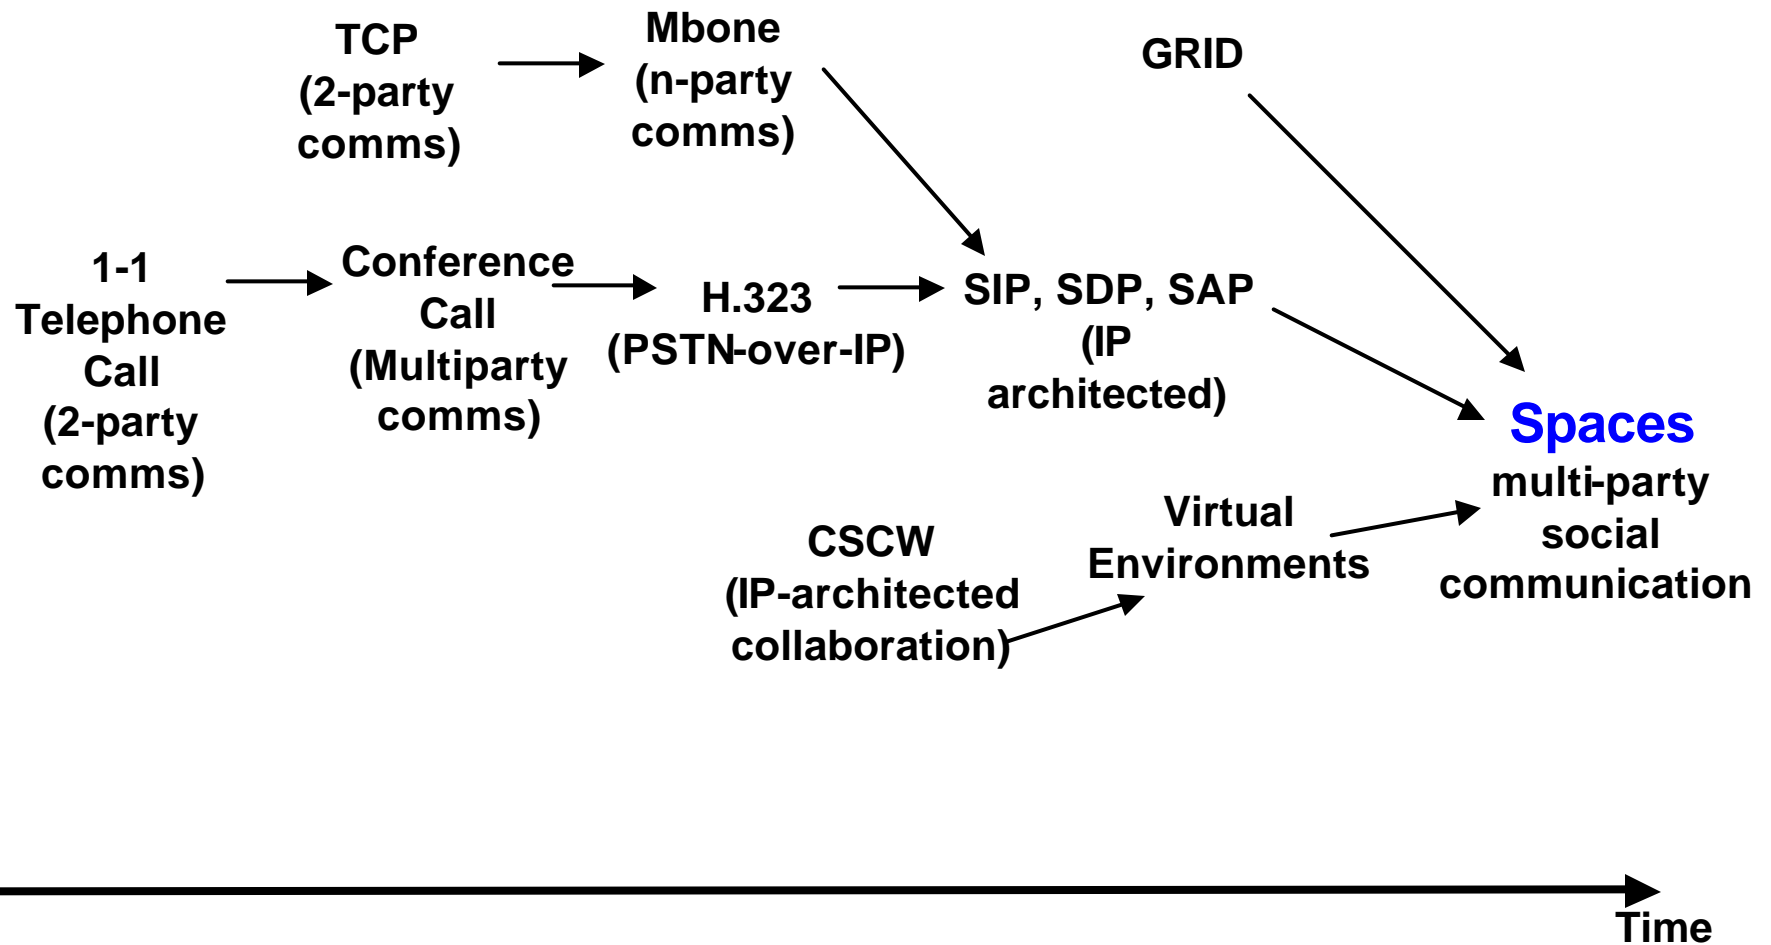

Architecture




Spaces

- Similar concerns to sessions
 - Group membership, invite and join protocols, security, characteristics of exchanged media
- No need to reinvent everything
 - Extended SIP,SDP, RTP interface to clients
- Rich Web Services interface to applications
 - Construction of world and policy management

Evolution of Communications & Session Technologies

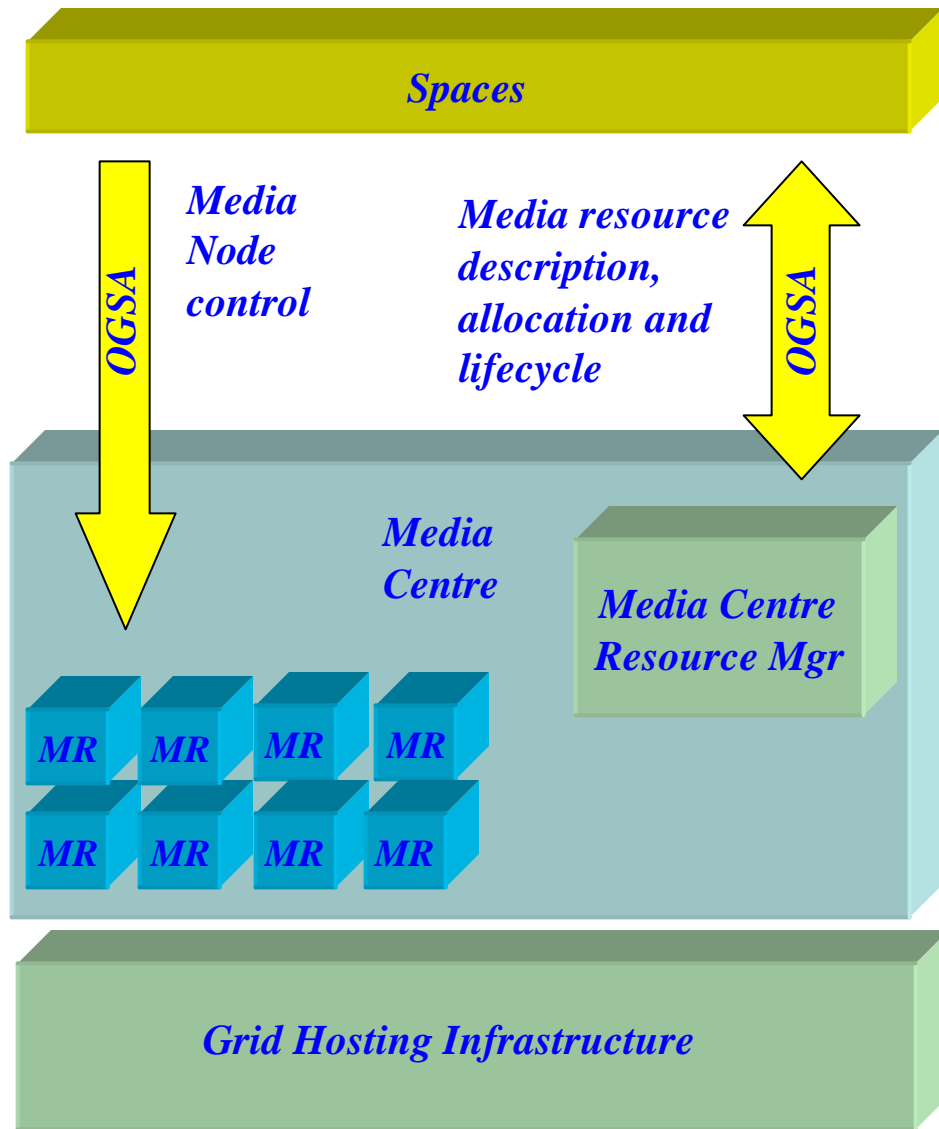


Media Resources




- Building blocks for media processing
- Palette of media functionality
 - Mixing, transcoding, rebroadcasting, record, playback
- Allocated on demand
- Interface exposed as OGSA

Media Grid Architecture




- Media Centre manages pool of Media Resources
 - Provides virtualisation
 - Implementation optimisation
- Sophisticated, fluid audio/media bridge
- OGSA as a control and management technology
- OGSA interface to describe, allocate and control mesh of resources and media interconnection relationships


Mixing architectures for Communication




a) Full mesh endpoint mixing



b) Centralised mixing



c) Multicast endpoint mixing



d) Hierarchical mixing

Why the GRID?

- Ability to manage distributed computing resources on a large scale
 - Service lifecycle management, Security
 - Dynamic allocation to meet demand
- Standard interfaces - OGSA
 - Resource virtualisation
 - OS independence, Soft vs. H/W solutions
 - Uniform Service abstraction for media resources
 - Quickly introduce new services
 - Event notification
 - Intuitive object model
- 21st Century software engineering and deployment technology for Communications
 - Compare to ASN.1 binary encoding and IN protocols

Summary

Problem

Emerging next generation communication applications are not well supported by current communication infrastructure

Drivers

Trend to large-scale sophisticated Virtual Meeting Places

Demand for computation, bandwidth and storage

Opportunity

Use GRID technology to support large-scale multi-party communication and real-time media processing

Challenge

Description, deployment, control, and management of mesh of interconnected media resources to form Media Grid